

IN THE CLAIMS:

- 1 1. (Currently Amended) A method for a storage operating system implemented in a
2 storage system to concurrently perform readahead operations for a plurality of different
3 read streams established in one or more files, directories, vdisks or luns stored in the
4 storage system, the method comprising:
5 allocating at least one readset data structure (“readset”) for each of the one or
6 more files, directories, vdisks or luns in which the plurality of different read streams is
7 established, wherein the number of readsets allocated for each file, directory, vdisk or lun
8 depends on the size of that file, directory, vdisk or lun;
9 receiving a client read request at the storage system, the client read request
10 indicating client-requested data for the storage operating system to retrieve from a file,
11 directory, vdisk or lun stored in the storage system;
12 determining whether the received client read request matches any of ~~a~~the
13 plurality of readset data structures (“readsets”) allocated for the file, directory, vdisk or
14 lun containing the client-requested data; and
15 performing readahead operations in accordance with a set of readahead metadata
16 stored in an associated readset that is determined to match the received client read
17 request, wherein the readahead metadata describes the associated readset.
- 1 2. (Currently Amended) The method of claim 1, further comprising:
2 ~~allocating at least one readset for each of the one or more files, directories, vdisks~~
3 ~~or luns in which the plurality of different read streams is established;~~
4 generating a separate set of readahead metadata for each of the plurality of
5 different read streams; and
6 storing each generated set of readahead metadata in a different readset allocated
7 for the file, directory, vdisk or lun in which the read stream associated with the generated
8 set of readahead metadata is established.
- 1 3. (Original) The method of claim 1, further comprising:

2 initializing each allocated readset to store a predetermined set of values.

1 4. (Cancelled)

1

1 | 5. (Currently Amended) The method of claim ~~4~~2, wherein the number of readsets
2 allocated for a file, directory, vdisk or lun is dynamically increased as the size of that file,
3 directory, vdisk or lun is increased.

1 6. (Original) The method of claim 1, wherein a first readset is determined to match the
2 received client read request if the first readset stores a set of readahead metadata
3 associated with a read stream that is extended by the client-requested data.

1 7. (Original) The method of claim 1, wherein a second readset is determined to match
2 the received client read request when the client-requested data is located within a
3 predetermined fuzzy range associated with the second readset.

1 8. (Original) The method of claim 7, wherein the fuzzy range is derived based on a
2 multiple of a number of client-requested data blocks specified in the received client read
3 request.

1 9. (Original) The method of claim 7, wherein the fuzzy range extends in both a forward
2 direction and a backward direction in relation to a last data block retrieved in a read
3 stream associated with the second readset.

1 10. (Original) The method of claim 1, wherein a third readset is determined to
2 match the received client read request if the third readset is determined to be unused.

1 11. (Original) The method of claim 10, wherein the third readset is determined to be
2 unused when a level value stored in the third readset equals a special indicator value.

1 12. (Original) The method of claim 1, wherein readahead operations are not performed if
2 the storage operating system determines that the file, directory, vdisk or lun containing
3 the client-requested data is accessed using a random access style.

1 13. (Original) The method of claim 12, wherein a DAFS cache hint included in
2 the received client read request indicates that the file, directory, vdisk or lun containing
3 the client-requested data is accessed using a random access style.

1 14. (Original) The method of claim 1, wherein readahead operations are not
2 performed unless:
3 (i) a readset is determined to match the received client read request; and
4 (ii) the matching readset stores a set of readahead metadata associated
5 with a read stream that is extended by the client-requested data past a
6 predetermined data block or memory address.

1 15. (Original) The method of claim 1, further comprising:
2 if the received client read request does not match any of the readsets allocated for
3 the file, directory, vdisk or lun containing the client-requested data, then performing the
4 steps:
5 identifying the received client read request as being the first read
6 request in a new read stream;
7 generating a set of readahead metadata associated with the new
8 read stream;
9 selecting for reuse one of the readsets allocated for the file,
10 directory, vdisk or lun containing the client-requested data; and
11 storing the generated set of readahead metadata associated with the
12 new read stream in the readset selected for reuse.

1 16. (Original) The method of claim 15, wherein the readset selected for reuse stores a
2 level value that is less than or equal to level values stored in each of the other readsets
3 associated with the file, directory, vdisk or lun containing the client-requested data.

1 17. (Original) The method of claim 1, wherein the client read request received at the
2 storage system is a file-based client read request.

1 18. (Original) The method of claim 1, wherein the client read request received at
2 the storage system is a block-based client read request.

1 19-28 (Cancelled)

1 29. (Currently Amended) A storage system that employs a storage operating system to
2 concurrently perform readahead operations for a plurality of different read streams
3 established in one or more files, directories, vdisks or luns stored in the storage system,
4 the method storage system comprising:
5 means for allocating at least one readset data structure ("readset") for each of the
6 one or more files, directories, vdisks or luns in which the plurality of different read
7 streams is established, wherein the number of readsets allocated for each file, directory,
8 vdisk or lun depends on the size of that file, directory, vdisk or lun;
9 means for receiving a client read request at the storage system, the client read
10 request indicating client-requested data for the storage operating system to retrieve from a
11 file, directory, vdisk or lun stored in the storage system;
12 means for determining whether the received client read request matches any of a
13 the plurality of readset data structures ("readsets") allocated for the file, directory, vdisk
14 or lun containing the client-requested data; and
15 means for performing readahead operations in accordance with a set of readahead
16 metadata stored in an associated readset that is determined to match the received client
17 read request, wherein the readahead metadata describes the associated readset.

1 30. (Currently Amended) A computer-readable media comprising instructions for
2 execution in a processor for the practice of a method for a storage operating system
3 implemented in a storage system to concurrently perform readahead operations for a
4 plurality of different read streams established in one or more files, directories, vdisks or
5 luns stored in the storage system, the method comprising:

6 allocating at least one readset data structure (“readset”) for each of the one or
7 more files, directories, vdisks or luns in which the plurality of different read streams is
8 established, wherein the number of readsets allocated for each file, directory, vdisk or lun
9 depends on the size of that file, directory, vdisk or lun;
10 receiving a client read request at the storage system, the client read request
11 indicating client-requested data for the storage operating system to retrieve from a file,
12 directory, vdisk or lun stored in the storage system;
13 determining whether the received client read request matches any of ~~a~~ the
14 ~~plurality of readset data structures (“readsets”)~~ allocated for the file, directory, vdisk or
15 lun containing the client-requested data; and
16 performing readahead operations in accordance with a set of readahead metadata
17 stored in an associated readset that is determined to match the received client read
18 request, wherein the readahead metadata describes the associated readset.

1 Please add new claims 31 *et al.*

1 31. (New) A method for a storage operating system implemented in a storage system to
2 concurrently perform readahead operations for a plurality of different read streams
3 established in one or more files stored in the storage system, comprising:
4 allocating at least one read set data structure (“readset”) for each of the one or
5 more files, directories, vdisks or luns in which the plurality of different read streams is
6 established wherein the number of readsets allocated for each file depends on the size of
7 that file;
8 generating a separate set of readahead metadata for each of the plurality of
9 different read streams; and
10 storing each generated set of readahead metadata in a different readset allocated
11 for the file in which the read stream associated with the generated set of readahead
12 metadata is established;
13 receiving a client read request at the storage system, the client read request
14 indicating client-requested data for the storage operating system to retrieve from a file,
15 stored in the storage system;
16 determining whether the received client read request matches any of a plurality of
17 readsets allocated for the file containing the client-requested data; and
18 performing readahead operations in accordance with a set of readahead metadata
19 stored in a readset that is determined to match the received client read request.

1 32. (New) The method of claim 31, wherein the file is broad term describing either a file,
2 directory, vdisk or lun.

1 33. (New) The method of claim 31, further comprising:
2 initializing each allocated readset to store a predetermined set of values.

1 34. (New) The method of claim 31, wherein the number of readsets allocated for a file is
2 dynamically increased as the size of that file is increased.

- 1 35. (New) The method of claim 31, wherein a first readset is determined to match the
2 received client read request if the first readset stores a set of readahead metadata
3 associated with a read stream that is extended by the client-requested data.
- 1 36. (New) The method of claim 31, wherein a second readset is determined to match the
2 received client read request when the client-requested data is located within a
3 predetermined fuzzy range associated with the second readset.
- 1 37. (New) The method of claim 36, wherein the fuzzy range is derived based on a
2 multiple of a number of client-requested data blocks specified in the received client read
3 request.
- 1 38. (New) The method of claim 36, wherein the fuzzy range extends in both a forward
2 direction and a backward direction in relation to a last data block retrieved in a read
3 stream associated with the second readset.
- 1 39. (New) The method of claim 31, wherein a third readset is determined to match the
2 received client read request if the third readset is determined to be unused.
- 1 40. (New) The method of claim 39, wherein the third readset is determined to be unused
2 when a level value stored in the third readset equals a special indicator value.
- 1 41. (New) The method of claim 31, wherein readahead operations are not performed if
2 the storage operating system determines that the file, directory, vdisk or lun containing
3 the client-requested data is accessed using a random access style.
- 1 42. (New) The method of claim 41, wherein a DAFS cache hint included in the received
2 client read request indicates that the file, directory, vdisk or lun containing the client-
3 requested data is accessed using a random access style.

1 43. (New) The method of claim 31, wherein readahead operations are not performed
2 unless:
3 (i) a readset is determined to match the received client read request; and
4 (ii) the matching readset stores a set of readahead metadata associated
5 with a read stream that is extended by the client-requested data past a
6 predetermined data block or memory address.

1 44. (New) The method of claim 31, further comprising:
2 if the received client read request does not match any of the readsets allocated for
3 the file, directory, vdisk or lun containing the client-requested data, then performing the
4 steps:
5 identifying the received client read request as being the first read
6 request in a new read stream;
7 generating a set of readahead metadata associated with the new
8 read stream;
9 selecting for reuse one of the readsets allocated for the file,
10 directory, vdisk or lun containing the client-requested data; and
11 storing the generated set of readahead metadata associated with the
12 new read stream in the readset selected for reuse.

1 45. (New) The method of claim 44, wherein the readset selected for reuse stores a level
2 value that is less than or equal to level values stored in each of the other readsets
3 associated with the file, directory, vdisk or lun containing the client-requested data.

1 46. (New) The method of claim 31, wherein the client read request received at the storage
2 system is a file-based client read request.

1 47. (New) The method of claim 31, wherein the client read request received at the storage
2 system is a block-based client read request.